

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau


INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 9/46		A1	(11) International Publication Number: WO 98/43163
			(43) International Publication Date: 1 October 1998 (01.10.98)
(21) International Application Number: PCT/US98/05779 (22) International Filing Date: 24 March 1998 (24.03.98) (30) Priority Data: 08/822,680 24 March 1997 (24.03.97) US (71) Applicant: EMC CORPORATION [US/US]; 171 South Street, Hopkinton, MA 01748-9103 (US). (72) Inventors: RAZ, Yoav; 15 Vine Street, Newton, MA 02167 (US). CARITI, Walter, Michael; 32R Hill Street, Medway, MA 02053 (US). JOHNSON, Douglas, V.; 146 Church Street, Marlborough, MA 01752 (US). (74) Agent: PRAHL, Eric, L.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).			(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published With international search report.

(54) Title: AN AGENT-IMPLEMENTED LOCKING MECHANISM

(57) Abstract

A method of synchronizing access to a data object that is stored within a shared storage area of a data storage system to which a plurality of host processors are connected, the method including the steps of establishing one of the plurality of host processors as a master host processor, wherein the master host processor owns the shared object within the data storage system; sending an access request from a slave host processor to the master host processor to access the shared object within the data storage system, wherein the slave host processor is one of the plurality of host processors other than the master host processor; in response to receiving the access request at the master host processor, checking whether any conflicting locks are still pending from prior access requests from any of the plurality of host processors; if no locks are still pending, granting the access request from the slave host processor; and notifying the slave processor whether its access request was granted.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

- 1 -

AN AGENT-IMPLEMENTED LOCKING MECHANISMBackground of the Invention

The invention relates generally to locking
5 mechanisms for controlling access to shared data storage
resources.

Today, data storage systems are commercially
available that provide many gigabytes of data storage
capacity and that permit multiple host processors to
10 connect to and access that storage. The Symmetrix Model
55XX series which is sold by EMC Corporation of
Hopkinton, MA is one example of such a system. The
Symmetrix is a high performance Integrated Cache Disk
Array (ICDA) that is designed for online storage. As is
15 typical of other of the commercially available designs,
the Symmetrix Model 55XX includes a large array of
smaller storage (e.g. 5-1/4 inch) disk devices and a
cache for temporarily storing data which is transferred
into and out of the data storage system. In the Model
20 55XX there can be as many as 128 such disk devices. Each
individual disk device is capable of storing from 2 to 9
GB. Thus, the total storage capacity of the system is in
the hundreds of Gigabytes.

In the Symmetrix Model 55XX ICDA's one is able to
25 define a portion of the data storage area that is shared
by two or more of the host processors that are connected
to the system. If the host processors that are able to
access files in the shared storage area are all running
UNIX operating systems, a problem can occur. Since one
30 Unix system does not know anything about what another
Unix system might be doing, it is very possible to
corrupt shared files. That is, recently updated data in
one file can be inadvertently overwritten by another host
processor that is connected to the data storage system.

- 2 -

The invention described below solves this problem, not only for a group of host processors running homogeneous operating systems (e.g. UNIX based systems) but also for any group of host processors that are
5 running heterogeneous operating systems.

Summary of the Invention

In general, in one aspect, the invention is a method of synchronizing access to a data object that is stored within a shared storage area of a data storage
10 system to which a plurality of host processors are connected. The method includes the steps of establishing one of the host processors as a master host processor, according to which it owns the shared object within the data storage system; sending an access request from a
15 slave host processor to the master host processor to access the shared object within the data storage system; in response to receiving the access request at the master host processor, checking whether any conflicting locks are still pending from prior access requests from any of
20 the host processors; if no locks are still pending, granting the access request from the slave host processor; and notifying the slave processor whether its access request was granted.

Preferred embodiments include the following
25 features. The method further includes the step of denying the access request from the slave host processor if a conflicting lock is still pending. Also, the step of granting an access request involves recording in local memory that the access request was granted to the
30 requesting processor. The object is a file system and the access request is a request to access a file in the file system.

In general, in another aspect, the invention is agent software on computer-readable medium. The agent

- 3 -

software is for running on a master host processor which is one of a plurality of host processors connected to a data storage system and the host processors all have access to a shared storage area in the data storage system. The agent software includes means for receiving an access request from another one of the host processors, wherein the access request is a request to access an object in the data storage system; means for determining whether any conflicting locks are pending for the object stored within the shared region of the data storage system; means for granting the access request from the requesting host processor if it is determined that no conflicting locks are pending for any other host processor; and means for notifying the requesting host processor whether its access request was granted.

In preferred embodiments, the agent software also includes means for denying the access request from the requesting processor if it is determined that a conflicting lock is pending for another host processor, and the notifying means notifies the requesting host processor whether its access request was granted or denied.

In general, in still another aspect, the invention is interface software on computer-readable medium. The interface software is for running on a slave host processor which is one of a plurality of host processors connected to a data storage system, all of which have access to a shared storage area in the data storage system. The interface software includes means for preventing the slave host processor from accessing an object stored in the shared storage area when no authorization has been received from a master host processor which is another one of the host processors; means for sending an access request from the slave host processor to the master host processor, the access

- 4 -

request being a request for accessing the object in the shared storage region; means for receiving an acknowledgment from the master host processor indicating whether the access request was granted; and means for
5 permitting the slave processor to access the object if the received acknowledgment grants the access request.

In preferred embodiments, the interface software also includes means for notifying the master host processor when an access operation involving the object
10 is complete.

In general, in yet another aspect, the invention is an access control method running on a host processor which is one of a plurality of host processors connected to a data storage system, all of which have access to a
15 shared storage area in the data storage system. The method includes the steps of receiving an access request from another one of the host processors, the access request being a request to access an object stored in the shared storage area of the data storage system;
20 determining whether any conflicting locks are pending for the object stored within the shared region of the data storage system; granting the access request from the requesting host processor if it is determined that no conflicting locks are pending for any other host
25 processor; and notifying the requesting host processor whether its access request was granted.

In preferred embodiments, the method also includes denying the access request from the requesting processor if it is determined that a conflicting lock is pending
30 for another host processor, and the step of notifying comprises notifying the requesting host processor whether its access request was granted or denied.

In general, in another aspect, the invention is an access control method running on a slave host processor
35 which is one of a plurality of host processors connected

- 5 -

to a data storage system, and all of which have access to a shared storage area in the data storage system. The method includes preventing the slave host processor from accessing an object stored in the shared storage area

5 when no authorization has been received from a master host processor which is another one of the host processors; sending an access request from the slave host processor to the master host processor, the access request being a request for accessing the object in the

10 shared storage region; receiving an acknowledgment from the master host processor indicating whether the access request was granted; and permitting the slave processor to access the object if the received acknowledgment grants the access request.

15 In preferred embodiments, the method also includes notifying the master host processor when an access operation involving the object is complete.

The invention is a mechanism for synchronizing accesses by multiple hosts to shared files on a data

20 storage system. The mechanism, which provides an asymmetrical solution to the synchronization problem, solves the problem with database sharing at the file level. Through an agent running on one of the host processors, it provides a locking functionality without

25 having to modify the operating system. The agent is simply an application, though in principle it could also be embedded in the operating system.

The invention also effectively and efficiently solves the inter-host access coordination problem when

30 MVS files in shared storage areas are being directly accessed from open systems through SCSI channels. The solution is general and allows multi users to access shared files in a multi host environment (a combination of MVS based and open system hosts).

- 6 -

Other advantages and features will become apparent from the following description of the preferred embodiment and from the claims.

Brief Description of the Drawings

5 Fig. 1 is a block diagram of a system which embodies the invention; and

 Fig. 2 is a flow chart of the process for accessing files in shared data storage area.

Description of the Preferred Embodiments

10 System Overview:

 In Fig. 1, there are several host processors 10(a), 10(b), ...10(n) (referred to generally as host processors 10) individually and separately connected to a common data storage system 12 through respective channels 15 13(a)-13(n). Data storage system 12 is configured to provide a shared storage region that can be accessed by all of the host computers 10 connected to the data storage system. The shared storage region can include all of the disk devices, some subset of the disk devices, 20 or part of a single disk device. In any case, each host processor 10 is able to access the data (e.g. files) that is stored within the shared storage region. In the described embodiment, data storage system 12 is a Symmetrix Model 55XX series integrated cache disk array 25 including an array of storage disk devices 14 and cache memory 16, though it could be any other data storage system that can be configured to provide shared storage area for multiple processors that are connected to it.

 There is a channel 18 through which the host 30 processors can communicate with each other. That channel might be implemented by a network or communication link that is separate from the data storage system as is

- 7 -

shown, or it might be implemented by the data storage system itself.

In the described embodiment, all of the host processors have a UNIX operating system, though not all of the UNIX operating systems need be the same. Indeed, it is anticipated that each of the host processors might have a different UNIX operating system, e.g. Sun/Solaris, AIX, Digital UNIX, HP-UX, AS-400, etc.

For the following description, it is assumed that host processor 10(a) has a different operating system from the other host processors and that it has created a file system within the shared region of data storage system 12. In other words, for the file system in data storage system 12, the operating system of host processor 10(a) is the native operating system, i.e., the file system is defined in accordance with the rules of that operating system.

Because the file system is implemented within the shared storage region, the other host processors are also able to access the files in the file system. The host processor 10(a) that created the file system is considered to be the "owner" of the file system and is generally referred to as the "master" host processor. The other host processors are referred to as "slave" host processors.

The master host processor 10(a) is different from the other host processors 10(b)-10(n) in that it also has an agent software program 20 (also referred to generally as agent 20) running on it; and the other host processors have interface software programs 22 (also referred to generally as interfaces 20) running on them.

Agent 20 synchronizes access by all host processors to the files within the file system in data storage system 12. It does this by determining which host processor is authorized to have access at any given

- 8 -

time and denying access to all other host processors. That is, it implements a locking function for accesses to the shared file system so as to prevent inadvertent corruption of the data in the file system. Agent 20
5 implements whatever locks are appropriate for achieving the type of access control that is desired. This might include, for example, exclusive and/or shared locks. It might also include other lock types such as those that are well known to persons skilled in the art.

10 The interface 22 running on the other host processors has several primary functions. One function is simply to communicate with the agent (e.g. over channel 18) and obtain the authorization that is required before accessing the file will be permitted. A second
15 related function is to prevent access to the file until specific authorization is given for such access. A third function is to handle the communications between its host processor and the file system. If the file system is owned by a host processor that has a different operating
20 system, then the interface is constructed to enable the requesting host processor to understand and manipulate the file structure of the other operating system. That is, in accordance with this last primary function interface 22 gives knowledge of the file system structure
25 to the non-native operating system on the slave host processor. Note, however, that if identical operating systems are running on all of the host processors, then no mapping of services needs to be performed to enable the other host processors to know how to understand and
30 manipulate the file system.

On each of the host processors there is also a management module 24 that provides management information which is necessary for the operation of the agent 20 and the interfaces 22. For example, the management module
35 identifies which host processor owns the file system.

- 9 -

That is, the management function identifies on which host processor the agent for a particular file system is located.

If different operating system are used on each of the other host processors, then the host processor that owns a particular file system is the host processor that is running the native operating system for the file system. On the other hand, if it is the case that all of the host processors are running the same operating system, then the assignment of file system ownership is completely arbitrary. That is, ownership can be assigned to any one of the host processors.

It may also be the case that there are multiple file systems, each having a different native operating system. For example, one host processor running operating system A might create a file system A and a different host processor running an operating system B might create a file system B, both file system in the shared storage region of data storage system 12. In that case, there would need to be two agents, one agent running on the host processor with operating system A and another agent running on the host processor with operating system B.

In any event, the management module is typically configured upon system set up. At that time, the relevant ownership information is stored into the local management module and the location of the agents is established. The management information can be made available from a central location to all of the hosts that are participating. When a new host processor is added to the system, it can simply look at management information to learn which processor owns what file system.

Actually, the management function can be either centralized or distributed. If it is distributed, some

- 10 -

communication mechanisms will be needed to distribute the table identifying ownership etc. to all hosts where it is can then be stored locally. In that case, the table could be updated whenever a host makes an open file request. That is, part of the process of opening the file could also check ownership information that is stored in the central location.

Access Synchronization Protocol:

The protocol which is implemented by the agent and the interface components operates as illustrated in Fig. 2.

When software running on a host processor seeks to access a file that is stored in the shared region of the data storage system, it generates a file access request (e.g. open file) which passes to the interface program on that host processor (step 100). The interface then processes the file access request by first determining, through referencing its associated management module 24, which host processor is the owner of the file to which access is being requested (step 102). After determining which host processor owns the relevant file system, the interface sends a request to open that file to the agent running on the identified host processor, i.e., the master host processor for that file system (step 104). In other words, the file access request does not go directly to the data storage system, as would normally be the case, but rather the interface processes it and redirects it to the appropriate agent to obtain authorization to access the file in the data storage system. Without authorization to do so, the interface will not permit the host processor to access the target file.

As noted above, the means by which the request is sent to the agent is not important. It could be sent

- 11 -

over a separate channel or it could be sent via the data storage system itself, such as is described in U.S.S.N. 08/616,486, incorporated herein by reference.

The agent on the master host processor receives
5 the request (step 106) and then determines whether any other host processors currently have a conflicting lock on the file (step 108). The agent makes this determination by referring to a locally maintained table
30 (see Fig. 1) in which the agent records the granting and the release of locks on a per file basis. It may be
10 that a prior lock was granted to another host processor but that the present file access request is not inconsistent with the lock that is still pending. For example, if a shared lock was granted and the new file
15 access request is only a read/no-write request, then allowing the file access will not be inconsistent with the pending lock and it will be authorized.

If no conflicting lock is still pending, then the agent is free to grant the appropriate authorization to
20 the requesting interface. Optionally, the agent may also open up the file on behalf of the requesting host processor or, alternatively, it may simply leave that responsibility to the interface in the requesting host processor. In either case, it sends back an
25 acknowledgment to the interface's request, authorizing access to the file (step 112) and it records the granting of that authorization in its local table.

When the interface at the requesting host processor receives the acknowledgment indicating that a
30 file access authorization (e.g. lock) has been given to it for accessing the file, it then carries out the operation that was originally requested by the application running on the host processor (step 114).

Upon completing the file operation, the interface
35 on the host processor sends a notification to the agent

- 12 -

releasing the lock (step 116). The agent then clears its table of the lock so that other host processors can be granted authorization to access the file (step 118).

In step 108, if the agent is unable to authorize
5 the requested file access operation because a conflicting lock has already been granted and is still pending, it sends a message to the interface denying its request (step 110). In that case, the interface notifies the application that generated the file access request and
10 waits for a new request to be generated at a later time.

Note that the request that the interface sends to the agent and the returned acknowledgment might typically include a set of parameters which are used to specify the particular file access operation that is being requested
15 and they type of lock that was requested and the type of lock that was granted.

In the above description, it had been assumed that the operating systems were all open systems type operating systems (i.e., UNIX). However, that need not
20 be the case. Indeed, the Symmetrix Model 55XX series ICDA's allow both mainframe and open systems devices to connect to the storage system at the same time and to access a shared storage region within the storage system. Thus, at least one of the host processors might be
25 running an MVS operating system and the others might be running open systems operating systems. In the case of an MVS file system in the data storage system, there would need to be an MVS based agent on the mainframe host processor. That agent would acquire and release locks on
30 behalf of the open system host processors that want to access any MVS files in the data storage system. That is, prior to accessing an MVS file by the open system in a certain locking mode (e.g., DISP=SHR), the agent attempts to open the file in the same mode. The open
35 system would access the MVS file in the requested mode if

- 13 -

and only if the agent succeeded to open the file in this mode. After processing of the file is complete, the lock is released by the agent closing the file. In this way the open system users accessing an MVS file acquires the
5 MVS lock by being an MVS user (via the agent).

As indicated previously, the interfaces on the open system side handle the job of communicating with the agent and getting authorization to access file. The interface also must be designed to understand the MVS
10 file system, i.e., it must be programmed to know the rules and protocol for the MVS file system.

Though the agent/interface mechanism was described in connection with a shared file system, it could be implemented in connection with any application or object
15 level sharing, e.g. storage level, file level, or database level.

Also, it should be understood that the functionality of the agent can be extended beyond that described above. For example, if the agent is running on
20 an MVS system, it could acquire MVS Catalog information about a requested file for the open systems (e.g., the VOLUME on which the file is located) in order to simplify the access to the file from the open systems.

In addition, if the same functionality is applied
25 to database files, the MVS agent would run as an application on the respective (MVS based) database system. The locks utilized could then be those available in the database system. If no explicit locking is available, implicit locking would need to be invoked by
30 the agent invoking appropriate database access operations (e.g. the minimum necessary for achieving the desired locking). For coarse granularity locking intended for relatively long database accesses, such a mechanism can be reasonable. For fine granularity and short accesses,
35 however, this mechanism could generate contention.

- 14 -

It should be noted that the agent/interface mechanism will work best in environments in which explicit locks are available, i.e., the locks are decoupled from the I/O (input/output).

5 The agent and interface software resides on
computer readable medium which is accessible by the host
processor on which that software is running. The
computer readable medium could be RAM, ROM, disk drives,
or any other form of electronic storage that is used for
10 software programs.

Other embodiments of the invention are within the following claims.

What is claimed is:

- 15 -

Claims:

1. A method of synchronizing access to a data object that is stored within a shared storage area of a data storage system to which a plurality of host
5 processors are connected, said method comprising:
 establishing one of said plurality of host processors as a master host processor, wherein said master host processor owns the shared object within the data storage system;
10 sending an access request from a slave host processor to the master host processor to access the shared object within the data storage system, said slave host processor being one of said plurality of host processors other than the master host processor;
15 in response to receiving the access request at the master host processor, checking whether any conflicting locks are still pending from prior access requests from any of the plurality of host processors;
 if no locks are still pending, granting the access
20 request from the slave host processor; and
 notifying the slave processor whether its access request was granted.
2. The method of claim 1 further comprising the step of denying the access request from the slave host
25 processor if a conflicting lock is still pending and wherein notifying step includes notifying the slave processor whether its access request was granted or denied.
3. The method of claim 2 wherein the step of
30 granting an access request comprises recording in local memory that the access request was granted to the requesting processor.

- 16 -

4. The method of claim 3 wherein the object is a file system.

5. The method of claim 4 wherein the access request is a request to access a file in the file system.

5 6. Agent software on computer-readable medium, said agent software for running on a master host processor which is one of a plurality of host processors connected to a data storage system, said plurality of host processors all having access to a shared storage
10 area in the data storage system, said agent software comprising:
 means for receiving an access request from another one of said plurality of host processors, said access request being a request to access an object in the data
15 storage system;
 means for determining whether any conflicting locks are pending for the object stored within the shared region of the data storage system;
 means for granting the access request from the
20 requesting host processor if it is determined that no conflicting locks are pending for any other host processor; and
 means for notifying the requesting host processor whether its access request was granted.

25 7. The agent software of claim 6 further comprising means for denying the access request from the requesting processor if it is determined that a conflicting lock is pending for another host processor, and wherein the notifying means notifies the requesting
30 host processor whether its access request was granted or denied.

- 17 -

8. The agent software of claim 7 wherein the object is a file system.

9. The agent software of claim 8 wherein the access request is a request to access a file in the file
5 system.

10. Interface software on computer-readable medium, said interface software for running on a slave host processor which is one of a plurality of host processors connected to a data storage system, said
10 plurality of host processors all having access to a shared storage area in the data storage system, said interface software comprising:

means for preventing said slave host processor from accessing an object stored in the shared storage
15 area when no authorization has been received from a master host processor which is another one of said plurality of host processors;

means for sending an access request from the slave host processor to the master host processor, said access
20 request being a request for accessing said object in the shared storage region;

means for receiving an acknowledgment from the master host processor indicating whether the access request was granted; and

25 means for permitting the slave processor to access the object if the received acknowledgment grants the access request.

11. The interface software of claim 10 further comprising means for notifying the master host processor
30 when an access operation involving the object is complete.

- 18 -

12. The interface software of claim 11 wherein the object is a file system.

13. The interface software of claim 12 wherein the access request is a request to access a file in the
5 file system.

14. An access control method running on a host processor which is one of a plurality of host processors connected to a data storage system, said plurality of host processors all having access to a shared storage
10 area in the data storage system, said access control method comprising:

receiving an access request from another one of said plurality of host processors, said access request being a request to access an object stored in the shared
15 storage area of the data storage system;

determining whether any conflicting locks are pending for the object stored within the shared region of the data storage system;

granting the access request from the requesting
20 host processor if it is determined that no conflicting locks are pending for any other host processor; and

notifying the requesting host processor whether its access request was granted.

15. The method of claim 14 further comprising
25 denying the access request from the requesting processor if it is determined that a conflicting lock is pending for another host processor, and wherein the step of notifying comprises notifying the requesting host processor whether its access request was granted or
30 denied.

- 19 -

16. The method of claim 15 wherein the object is a file system.

17. The method of claim 16 wherein the access request is a request to access a file in the file system.

5 18. An access control method running on a slave host processor which is one of a plurality of host processors connected to a data storage system, said plurality of host processors all having access to a shared storage area in the data storage system, said
10 method comprising:
 preventing said slave host processor from accessing an object stored in the shared storage area when no authorization has been received from a master host processor which is another one of said plurality of
15 host processors;
 sending an access request from the slave host processor to the master host processor, said access request being a request for accessing said object in the shared storage region;
20 receiving an acknowledgment from the master host processor indicating whether the access request was granted; and
 permitting the slave processor to access the object if the received acknowledgment grants the access
25 request.


19. The method of claim 18 further comprising notifying the master host processor when an access operation involving the object is complete.

20. The method of claim 19 wherein the object is
30 a file system.


- 20 -

21. The method of claim 20 wherein the access request is a request to access a file in the file system.

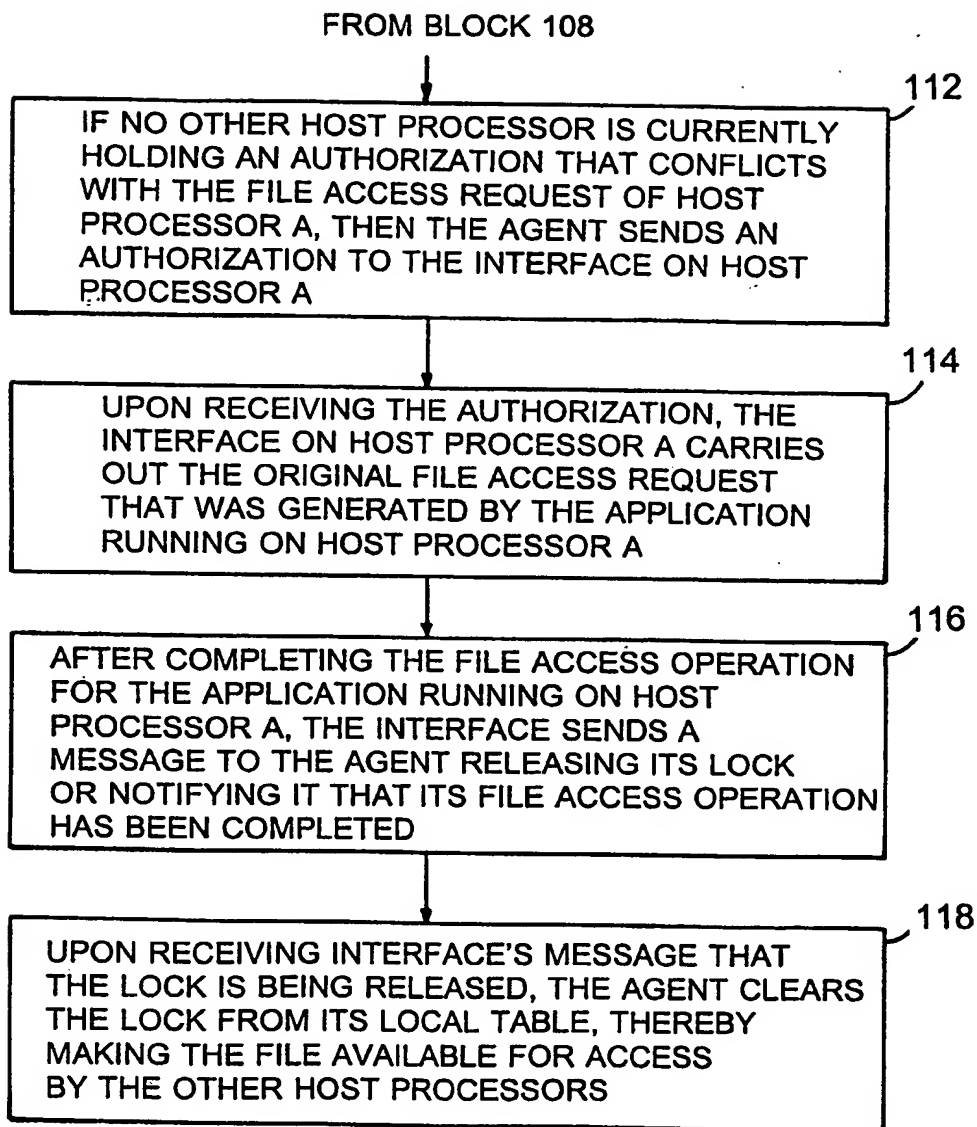
1/3

**FIG. 1**

2/3

**FIG. 2A**

3/3

**FIG. 2B**

INTERNATIONAL SEARCH REPORT

International Application No.
PCT/US 98/05779

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	GB 2 188 177 A (HITACHI LTD; HITACHI SOFTWARE ENG) 23 September 1987 see page 1, line 1 - page 4, line 23 see figure 1	1-21
X	WO 95 13583 A (CONNER KENNETH H) 18 May 1995 see page 4, line 25 - page 5, line 30 see page 11, line 1 - page 12, line 25 see figure 3	1-21



Further documents are listed in the continuation of box C.



Patent family members are listed in annex.

*** Special categories of cited documents :**

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "Z" document member of the same patent family

Date of the actual completion of the international search

30 June 1998

Date of mailing of the international search report

08/07/1998

Name and mailing address of the ISA
European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Brandt, J

INTERNATIONAL SEARCH REPORT

information on patent family members

International Application No.

PCT/US 98/05779

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
GB 2188177 A	23-09-1987	JP 62197858 A	01-09-1987
WO 9513583 A	18-05-1995	AU 1091295 A	29-05-1995